Heads-Up Face-Off:
On Style and Skill in the Game of Poker

Kevin Burns

The MITRE Corporation
202 Burlington Road
Bedford, MA 01730-1420 USA
kburns@mitre.org

Abstract

Poker is a game of imperfect information where players make inferences and investments under uncertainty. Winning requires skill as well as luck, and people play with different styles. Here I examine the roles of style and skill in a simplified poker game. I propose that style arises from mental limits in estimating odds and stakes, and I define style as a strategy for making decisions in light of these mental limits. I then develop formal models of cognitive styles, such as "Tight" and "Loose", and pit them against each other in heads-up (pair-wise) face-offs (match-ups). I also develop a formal model of normative skill and pit this skill against the styles. My results show that the best style for one player depends on the style of his opponent, and that Novice styles can be remarkably effective against Expert skill. I discuss the reasons behind these results and suggest how this study may apply to practical problems in business and warfare – which are real life games of inference and investment that are also played with style and skill.

Introduction

Poker is a game of *inference* and *investment* played with imperfect information. Players must infer the strength of their own cards relative to their opponents' cards based on the information they get from their opponents' bets. Players must then invest their chips based on the inferred strength of their cards. Thus, inferences affect investments, which in turn affect other players' inferences and investments in a complex cycle. The cycle is further complicated by the need to anticipate future game states as cards are dealt or drawn to each hand, and by the need to adjust inference and investment strategies as hands are won or lost and new hands are dealt and played.

The same complex cycle of inference and investment is at the crux of business, warfare and virtually every other domain of human endeavor (McDonald 1950). This makes poker a useful testbed for research on cognitive strategies, and in particular for exploring the roles of *style* and *skill* in mind games. Here I develop formal models of style and skill and pit them against each other in heads-up (pair-wise) face-offs (match-ups).

In general, style refers to *how* people play a game and skill refers to *how well* people play a game. The former (how) clearly affects the latter (how well), but the relation is not always obvious and in fact the colloquial distinction between style and skill is typically as follows: skill refers to aspects of a player's manner that are clearly linked to a desired outcome, while style refers to aspects of a player's manner that are not so clearly linked to an outcome. For example, in baseball, a "quick bat" would be considered batting skill while a "closed stance" would be considered batting style. Here my goal is to shed some light on the link between manner and outcome in the game of poker.

Taken to the extreme, one might argue that all style can be reduced to skill because all manner results in some outcome. But as a practical matter there will always be mental limits on what players can and will see as links between manner and outcome (often clouded by luck). Thus, people will always play with style as well as skill, and research on mental limits is key to understanding why people adopt different styles and how these styles evolve into skills – a process that leads to new styles at higher levels of skill.

Here I take a small step towards clarifying the relation between style and skill by developing formal models of cognitive strategies in a simplified poker game. I examine several Novice styles in head-to-head competition to see which ones are better than others (and why). I also examine how well the Novice styles perform against Extreme (bounding) styles and against Expert skill. The face-offs between Novice and Expert are particularly interesting because they illustrate the effectiveness of cognitive (heuristic) style against normative (optimal) skill.

Style and Skill

Superficially, poker is a game played with cards and chips where players use cards in their hands to win chips from the pot. Theoretically, poker is a game of inference and investment where a player must infer the strength of his own hand relative to his opponents' hands and then invest the appropriate chips. Mathematically, poker is a game of card probabilities and chip utilities.

Computationally, poker playing (Sklansky 1987) and many other practical problems can be boiled down to the math of odds and stakes. That is, a player's win:loss *odds* and cost:pot *stakes* determine whether he should "hold'em or fold'em" (as the song goes) in any betting situation.

For example, consider the simplest case where there are only two players, Player A and Player B, which in poker jargon is called a *heads-up* game (see Figure 3). Assume that both players have anted one chip to the pot and Player A has made a bet of two more chips (so pot=4 chips). Now Player B must make a choice. First, should he hold or fold? Then, if he decides to hold, should he call or raise?

Consider Player B's first choice: to hold or fold? As far as skill is concerned, the decision depends on his win:loss odds and the cost:pot stakes. For example, assume that Player B thinks his win:loss odds are 2:1, i.e., there is 2/3=67% chance that he (Player B) has the better hand and a 1/3=33% chance that his opponent (Player A) has the better hand.

When it is Player B's turn to act, the pot contains 4 chips and the cost to call the bet is 2 chips, so the cost:pot stakes are 2:4 (a call will make the pot 4+2=6 chips). The expected utility of a call by Player B is 6*(2/3)+0*(1/3)=4 chips, since a win will pay him 6 chips and a loss will pay him 0 chips. The expected utility of a fold by Player B is 2 chips, because he gives up the pot (with probability 1) but he retains (with probability 1) the 2 chips that he would have invested in a call. Since the expected utility of a call is 4 chips and the expected utility of a fold is only 2 chips, Player B should call (not fold).

As a simple heuristic for making this choice, Player B can compare his win:loss odds to the cost:pot stakes (sometimes called "pot odds", see Sklansky 1987). In the above example, the win:loss odds of 2:1 are better than the cost:pot stakes of 2:4=1:2 so Player B should call (not fold).

But now assume that, instead of 2:1, Player B thinks his win:loss odds are 1:2. That is, there is 1/3=33% chance that he (Player B) has the better hand and a 2/3=67% chance that his opponent (Player A) has the better hand. Since the win:loss odds are now the same as the cost:pot stakes, Player B has no rational basis for making a choice, i.e., his expected utility is the same for each option (call or fold). In this case the decision must be based on *style* rather than *skill*. Or, to be more formal:

Definition of Style: Style is a functional basis for making decisions in the absence of a rational basis for making decisions, i.e., when the expected utility for each option is the same.

On the surface it may seem like this definition leaves little room for style, since style will only apply when the expected utilities of various options are exactly equal. But this is true only if a player computes odds and stakes with *precision*, and then only if a player can make perfect *projections*. With respect to precision, people usually can and do estimate odds and stakes only approximately in their heads. For example, even though Player B knows for sure what hand he holds, he is unsure of what hand Player A holds so he may estimate his win:loss odds at a round number like 1:2 for a wide range of situations. With respect to projections, the above example was simplified in that Player B's action would end the hand. In other cases, e.g., if Player B chose to raise instead of call, then Player A might re-raise and Player B would have to consider this possibility (as well as his own possible responses to the raise) in order to estimate the projected odds and stakes. The projected win:loss odds will change because Player A's raise will give Player B more information about Player A's hand, and the projected cost:pot stakes will change because the raises will add more chips to the pot.

Referring to Figure 1, we can define a gray region where all choices are *rationally equivalent* because mental precision is not sufficient and mental projections are too difficult. It is in this gray region that choices are based on style rather than skill.

On the left and right of Figure 1, the two polar styles can be characterized as pessimistic and optimistic, respectively. A pessimistic style assumes that the gray region has negative expected utility (black), and an optimistic style assumes that the gray region has positive expected utility (white).
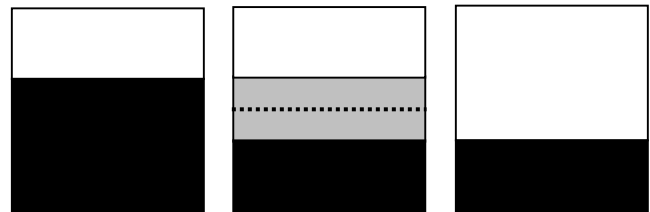
Figure 1. Gray shows region of rational equivalence in which choice is governed by style rather than skill. One style is pessimistic (left), where gray is assumed to be black (negative). The other is optimistic (right), where gray is assumed to be white (positive).

Tight and Loose

Returning to the above example, consider the choice faced by Player A when he made his original bet (after which Player B was forced to either raise, call or fold). Player A, who has to bet or fold at the start (see Figure 3), must consider the possibility that if he bets then Player B might raise. In making this choice Player A is faced with many questions, such as: What are the odds that his opponent holds a better hand? What is the chance that his opponent will raise, call or fold? How will he himself respond if his opponent raises the bet? What are the projected stakes, i.e., what will be the eventual size of the pot and how much will it cost to stay in the hand through all the raising and calling?

Unlike the call/fold choice faced by Player B, the bet/fold choice faced by Player A is riddled with problems of both precision and projection in estimating odds and stakes – i.e., the gray region (see Figure 1) is larger. In some cases Player A may hold a hand that is clearly strong (or weak) enough to justify a rational decision to bet (or fold). But in many cases his hand will be in the gray region so his decision will be based on style rather than skill. In poker jargon (Schoonmaker 2000), the two polar styles that apply to the bet/fold choice faced by Player A are *Tight* (pessimistic) and *Loose* (optimistic). That is, assuming a marginal hand (gray region of Figure 1), a Tight Player A would fold his hand while a Loose Player A would bet his hand.

Passive and Aggressive

To continue the example, assume that Player A makes a bet. Player B now faces two choices (see Figure 3). His first choice is between hold or fold; then, if he decides to hold, his next choice is between call or raise. The minimum hold is call, so the hold/fold choice is the same as the call/fold choice for Player B discussed earlier. Note that, unlike the bet/fold choice for Player A (above), a call or fold by Player B ends the hand (see Figure 3) so the cost:pot stakes can be estimated without making any projections, but the win:loss odds are still subject to mental limits of precision.

Assuming that Player B chooses to hold (not fold), he is then faced with the choice between raise or call. This choice involves problems of both precision and projection in estimating odds and stakes, much like those involved in the bet/fold choice of Player A. That is, Player B must consider what Player A will think of the raise and what Player A will do in response to the raise (at node A' of Figure 3). In many cases Player B's hand will not be clearly strong enough to raise so his choice will depend on style rather than skill. Here the poker jargon (Schoonmaker 2000) defines the two polar styles as *Passive* (pessimistic) and *Aggressive* (optimistic). That is, assuming a marginal hand (gray region of Figure 1), a Passive Player B would call the bet while an Aggressive Player B would raise the bet.

Four Style Profiles

The above examples illustrate two dimensions of a player's style that differ primarily in the stakes involved. That is, the Tight-Loose dimension of style applies to *Low Stakes* – i.e., making or calling a *bet*; while the Passive-Aggressive dimension of style applies to *High Stakes* – i.e., making or calling a *raise*. Based on these distinctions, poker analysts (Schoonmaker 2000) use a grid to define four basic styles of play (Figure 2): (T-A) Tight-Aggressive, (L-A) Loose-Aggressive, (L-P) Loose-Passive and (T-P) Tight-Passive. Some people define more or less styles, but these four seem to strike a good balance between simplicity and accuracy in the minds of most poker players. For example, poker champion Phil Hellmuth (2003) characterizes possible opponents using the following five "animal" (personality) types: *Lion* (T-A), *Jackal* (L-A), *Elephant* (L-P), *Mouse* (T-P) and *Eagle* (top pro). The Eagle is a special player who plays at a higher level where matters of style become matters of skill.

In this paper I model Hellmuth's five style profiles (Lion, Jackal, Elephant, Mouse and Eagle) at two skill levels: *Novice* (Lion, Jackal, Elephant, Mouse) and *Expert* (Eagle). In my models, the move from Novice to Expert involves replacing aspects of style by aspects of skill, to improve precision and projection in estimating win:loss odds and cost:pot stakes.
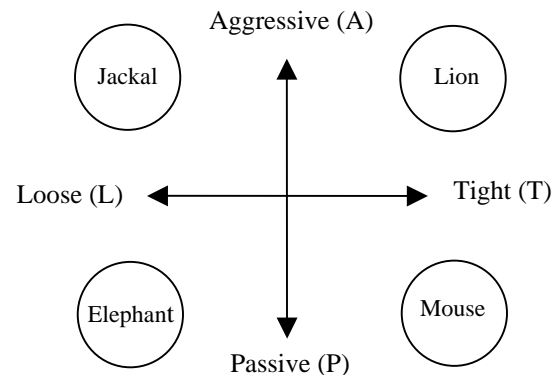

Figure 2. Four styles of poker playing.

Simplified Poker

Poker Versions

Various versions of poker differ primarily in the way cards are dealt or drawn. For example: In Draw Poker, all cards are dealt face down and some cards may be exchanged via discarding and drawing. In Stud Poker, some cards are dealt face up to each player. In Texas Hold'em, which has recently been popularized through media coverage of the "World Series of Poker" (McManus 2003), some cards are dealt face up and shared by all players.

Despite their many nuances, most standard poker games are similar in the way hands are ranked and the way chips are bet. With respect to ranking hands, there are ten classes of 5-card hands, e.g., Royal Flush, Straight Flush, Four-of-a-Kind, Full House, etc. These ranks are based on the relative rareness of the hand, but then there are further sub-ranks within each rank that are not based on rareness, e.g., Four Queens beats Four Jacks even though Queens are not more rare than Jacks.

With respect to betting chips, there is usually a pre-deal payment of chips to the pot as an *ante* (by all players) or a *blind bet* (by only some players, rotating with the deal). Then there are one or more rounds of betting, where the game state changes via dealing or drawing cards between each betting round. Within a given betting round, the betting rules are as follows: starting with one player (usually left of the dealer), the player can either *check or bet*. [A check is like a "pass" in which no chips are added to the pot but the player stays in the hand. In some games, checking is not allowed and the player is forced to *bet or fold*.] Once one player bets, then the options for subsequent players are *fold, call or raise*. [There is a limit on the number of raises per round, usually three.]

From a research perspective, if not a player's perspective, the problem with standard poker is that the games are intractable to analytical solution (Epstein 1977). In fact, even computer solutions are at or beyond the current state of the art (Billings et al. 2003; Koller and Pfeffer 1997) because, unlike chess and other games of perfect information, poker is a game of *imperfect information* where there are many possible game states to consider at each move (i.e., with a 52-card deck there are over 2.5 million 5-card combinations).

The cognitive (human) solution to this problem is to develop *satisficing strategies* (Simon 1981) that can be implemented by the unaided mind. Interestingly, these satisficing strategies are surprisingly successful in poker. In fact, computer researchers have found it difficult to develop software agents that can compete with champion players because the secret to winning poker lies not in exhaustively assessing all of the possibilities but rather in effectively assessing how your opponent will assess the possibilities. In short, the basic problem in skillful poker is *opponent modeling* (Billings et al. 2002).

In this paper I focus on a game of *Pared-down Poker* called *One Card High*. This game highlights the problem of opponent modeling and downplays the problem of enumerating possible 5-card hands that might be dealt or drawn from a 52-card deck. The game (see below) is played with a deck of only 11 cards. A similar but even simpler game, which used a 3-card deck, was studied by early game theorists (Nash and Shapley 1950; Kuhn 1950). The difference is that they had a mathematical interest in perfect players who could make infinite projections with infinite precision, while I have a psychological interest in human players who are bound by mental limits and hence must play with style.

Pared-down Poker

Goal: The object of the game is to win chips from a pot. The pot is won in a showdown when you have the best hand, or without a showdown when all opponents fold their hands.

Deck: The game is played with a deck of 11 cards numbered 10, 9, 8, 7, 6, 5, 4, 3, 2, Ace (1) and Joker (0). Each player is dealt one card.

Ranks: Each card is ranked according to its number. For example, 10 beats all other cards; 9 beats 8 and lower, etc. The Joker (0) is the worst card.

Ante: Before the deal, each player antes "a" chips to the pot.

Bets: The bet size "b" and raise size "r" can be "fixed limit", "spread limit", "pot limit" or "no limit".

Rounds: The game can involve several stages of dealing and betting. For example, a two-round game would involve ante, dealing (one card to each player), betting, dealing (a second card to each player) and betting. Players use the highest card in their hand to determine the winner in a showdown.

Game: Here I focus on the simplest version of *Pared-down Poker*, illustrated in Figure 3. This game, called *One Card High*, is played *heads-up* (two players) in one round with fixed limits of *ante=1 chip, bet=2 chips* and *raise=2 bets*. There is *no checking* allowed and only *one raise* allowed.
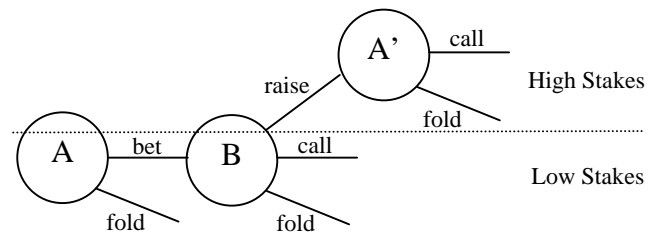
Figure 3. Game tree for Pared-down Poker (One Card High), played heads-up in one round with no checking and one raise. Player A must bet or fold. If Player A bets then Player B must raise, call or fold. If Player B raises then Player A (denoted A') must call or fold. The Tight-Loose dimension of style applies to Low Stakes decisions (below the dotted line) and the Passive-Aggressive dimension of style applies to High Stakes decisions (above the dotted line).

Modeling Players

Here I develop formal models of Hellmuth's (2003) five poker personalities, which are a mix of style and skill. There are four styles (Lion, Jackal, Elephant and Mouse) at Novice skill and one style (Eagle) at Expert skill.

The models are developed for the game of *One Card High* (see above) in which players of each style/skill play against one another *heads-up* (two players at a time) in a *face-off* (series of hands). In each face-off, all possible combinations of two one-card hands are dealt (11*10=110 deals are possible). Each player plays the role of A and B (see Figure 3) for each possible deal, making 220 games per face-off. In this way, for each possible deal, the face-off pits the Tight-Loose (Low Stakes) styles of two opponents against each other at nodes A and B of Figure 3 and the Passive-Aggressive (High Stakes) styles of the same two opponents against each other at nodes B and A'.

Novice Players: Lion, Jackal, Elephant and Mouse

At the Novice level, betting choices are based on the card(s) a player holds in his own hand (Helmuth 2003; Sklansky 1987) with little regard for what other players do. Thus, in the game of *One Card High*, the four basic styles (Figure 1) differ in how high a card a player needs in order to act at Low Stakes (Tight-Loose) or High Stakes (Passive-Aggressive).

Defining reasonable (not extreme) variations between styles (Figure 4), I assume that Loose will bet (or call a bet) with a 4 or higher card while Tight will bet (or call a bet) with a 6 or higher card. I also assume that Aggressive will raise (or call a raise) with a slightly better (+1) card, while Passive will require an even better (+2) card to raise. Thus, Tight-Passive will raise (or call a raise) with a 6+2=8; Tight-Aggressive will raise (or call a raise) with a 6+1=7; Loose-Passive will raise (or call a raise) with a 4+2=6; Loose-Aggressive will raise (or call a raise) with a 4+1=5 (see Figure 4).
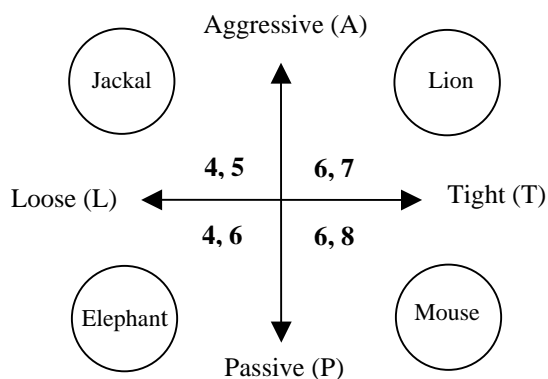
Figure 4. Minimum cards (X, Y) at which each style makes or calls a bet (X) and makes or calls a raise (Y).

At the Novice level, I assume that a player uses these (X, Y) card thresholds as the sole basis for all betting choices. I assume this even for the call/fold choices that will end the hand, where it is relatively easy to project cost:pot stakes, because estimates of the win:loss odds are still subject to metal limits. I also assume that Novice players ignore the information they get from their opponent's betting behavior, i.e., they do not update their win:loss odds.

In short, I assume that Novice players have (X, Y) thresholds for (Low, High) stakes. They then use X as the threshold criterion for Low Stakes choices (bet/fold at node A; call/fold at node B). Similarly, they use Y as the threshold criterion for High Stakes choices (raise/call at node B; call/fold at node A').

Expert Player: Eagle

The Expert player (Eagle) is different from the Novice players (Lion, Jackal, Elephant and Mouse) in three ways. First, he updates his win:loss odds based on the betting actions of his opponent. Second, he projects cost:pot stakes (and win:loss odds) for all possible actions that his opponent may take (Figure 3). Finally, he uses the updated/projected win:loss odds and cost:pot stakes to compute his expected utility at each choice point in Figure 3: bet or fold at node A; raise, call or fold at node B; call or fold at node A'.

The following two sections outline the Eagle's strategies for *updating odds* and *projecting stakes*, respectively.

Updating Odds

At the start of the game (node A in Figure 3), the Eagle Player A knows nothing about his opponent Player B's card except that his opponent's card cannot be the same as his own card (because there is only one of each card 0-10 in the deck). Thus, given a card Z in his own hand, the Eagle Player A estimates his win:loss odds as Z:(10-Z). For example, when Z=8 then his win:loss odds are 8:2=4:1.

At node B, an Eagle Player B knows more about his opponent Player A's card because Player A has made a bet (at node A). Here I assume the Eagle knows the style of his opponent and that he uses this information to update his own win:loss odds. For example, if the opponent Player A is a Mouse then Eagle Player B knows that his opponent holds a 6 or higher card. Thus, the Eagle can update his own win:loss odds from Z:(10-Z) to (Z-6):(10-Z). For example, if the Eagle holds an 8 then his win:loss odds after a Mouse has bet are (8-6):(10-8)=2:2 or 1:1.

Similarly, an Eagle Player A will update his win:loss odds at node A' after his opponent Player B has raised (see Figure 3). If Player B raises then it means he has a Y or better card, so an Eagle Player A will use the value of Y for his opponent's style (see Figure 4) to update his own win:loss odds for the card Z that he holds himself (like in the above example where Z=8 and X=6).

Projecting Stakes

Here I develop expressions for expected utility that the Eagle uses to make his choices: bet/fold at node A; raise/call/fold at node B; call/fold at node A'. I analyze the cases in reverse order (A', B then A, see Figure 3), moving from easier to harder projections.

At node A', the expected utility (U) for *call* is $U_{A',call} = P_{A',win}*(4b+2a)+(1-P_{A',win})*0$ chips, because the called pot will contain 4 bets (two from each player) and 2 antes (one from each player). The expected utility for *fold* is $U_{A',fold} = b$, because a player who folds gives up the pot but retains the b chips he would have spent (bet) to call at node A'. Thus, Player A should call rather than fold when $U_{A',call} > U_{A',fold}$ or $P_{A',win} > b/(4b+2a)$. Assuming b=2 and a=1, Player A should call (not fold) when $P_{A',win} > 1/5$, which is win:loss odds of 1:4 or better. Notice that, assuming baseline odds, Player A should call (not fold) at node A' if his card is a 2 (win:loss=2:8=1:4) or higher. However, because he uses updated odds, an Eagle (Expert) Player A will fold with a higher card. For example, if his opponent Player B is a Mouse then the Eagle will only call the raise if he holds a 9 or 10 himself.

At node B a player has three options: raise, call or fold. The expected utility for *raise* is $U_{B,raise} = P_{A',call|bet}*[P_{B,win|A',call}*(4b+2a)] + (1-P_{A',call|bet})*(3b+2a)$ chips. Here, $P_{A',call|bet}$ is the probability that Player A will call at node A' given that he has already bet at node A, and $P_{B,win|A',call}$ is the probability that B will win the pot given that A calls his raise. Note that $P_{B,win|A',call} \leq P_{B,win|A,bet}$. The first term in the expression for $U_{B,raise}$ arises because there will be 4b+2a chips in the pot (which Player B will win with probability $P_{B,win|A',call}$) if Player A calls at node A'. The second term arises because if Player A folds at node A' then Player B will win 3b+2a chips with probability 1.

The expected utility for *call* is $U_{B,call} = b+P_{B,win|A,bet}*(2b+2a)$. The first term arises because Player B will retain (with probability 1) the b chips that he would have spent to raise, and the second term arises because the called pot will contain 2b+2a chips (which Player B will win with probability $P_{B,win|A,bet}$). The expected utility for *fold* is 2b, since Player B will then retain (with probability 1) the two bets that it would cost him to raise.

Now the problem for Player B is that he must know $P_{A',call|bet}$ in order to determine which expected utility ($U_{B,raise}$ or $U_{B,call}$ or $U_{B,fold}$) is highest. Here I assume that, like the case of updating odds (see above), the Eagle knows his opponent's style and uses it to compute $P_{A',call|bet}$.

Finally, at node A the problem of projection is even more complex than at node B. In this case, an Eagle Player A must account for the three possible actions that his opponent Player B might take (i.e., raise, call or fold). Player A must also project his own win:loss odds, conditional on each of these actions by his opponent, plus project the cost:pot stakes and project what he himself will do at node A' in response to a raise by his opponent.

To start, the Eagle Player A considers what he will eventually do at node A'. If B raises then A knows that he himself will update (reduce) his own $P_{A,win}$ to $P_{A',win|B,raise}$ based on his knowledge of Player B's minimum raise card. Eagle Player A also knows that he himself will call at node A' (i.e., $P_{A',call|B,raise}=1$) if $P_{A',win|B,raise} > b/(4b+2a)$, otherwise he will fold (i.e., $P_{A',call|B,raise}=0$). Finally, Player A can use his knowledge of Player B's style (X, Y) to compute the probability that B will raise, call or fold at node B: $P_{B,raise}$, $P_{B,call}$ and $P_{B,fold}$. Then, the expected utility for *bet* by Player A at node A is as follows: $U_{A,bet} = P_{B,fold}*[b+(b+2a)] + P_{B,call}*[b+P_{A,win|B,call}*(2b+2a)] + P_{B,raise}*[P_{A',call|B,raise}* P_{A',win|B,raise}* (4b+2a) + P_{A',fold|B,raise}* (b)]$. This is compared to the expected utility for *fold* at node A, which is $U_{A,fold}=2b$.

Computed Results

Computer *face-offs* were held between each of the Novice styles (Lion, Jackal, Elephant, Mouse) in *heads-up* (pair-wise) competition. Each face-off was a set of 220 games that captured all 110 combinations of cards for the two different players with each combination played in each order (i.e., each player in role A and role B, see Figure 3). The result is an average net chips/game won by one style and lost by the other style (zero sum), as reported in Figure 5. In Figure 5, the arrowheads point to the winner of each pair-wise face-off and the numbers (and line thickness) reflect the magnitude of earnings by the winner.

The results show that one style, the Mouse, beats all other styles. However, the results are non-linear in that the Lion is more effective than the Mouse against the Jackal, even though the Mouse beats the Lion. The calculations were repeated using different (but not extreme, see below) card thresholds, with $X_{Loose} < X_{Tight}$ and $(Y-X)_{Aggressive} < (Y-X)_{Passive}$. The same pattern of results was obtained, namely: Tight is better than Loose; and Passive is better than Aggressive (for Tight or Loose).
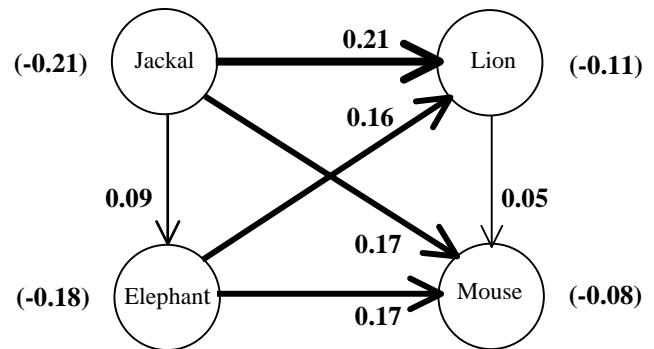
Figure 5. Results of pair-wise face-offs between each Novice style. Arrowheads point to winning style. Numbers are earnings (chips/game). Line thickness reflects earnings. Parenthetic numbers are earnings (negative) against the Eagle (Expert).

Further insight into these results was obtained by running more face-offs against *Extreme* styles. For these cases, I defined the following new styles: *Folder* (always folds – an Extreme Mouse), *Bettor* (always bets or calls at Low Stakes but never raises to or calls at High Stakes – an Extreme Elephant), *Raiser* (always bets and raises – an Extreme Jackal) and *Cheater* (a player who always knows his opponent's card and only bets or raises when he has a higher card himself). The minimum (X, Y) cards for these Extreme styles are: Folder (11, 11), Bettor (0, 11), Raiser (0, 0) and Cheater (Z+1, Z+1) where Z is the opponent's actual card.

Figure 6 shows the face-off results for the four Novice styles (Lion, Jackal, Elephant, Mouse) against the four Extreme styles.
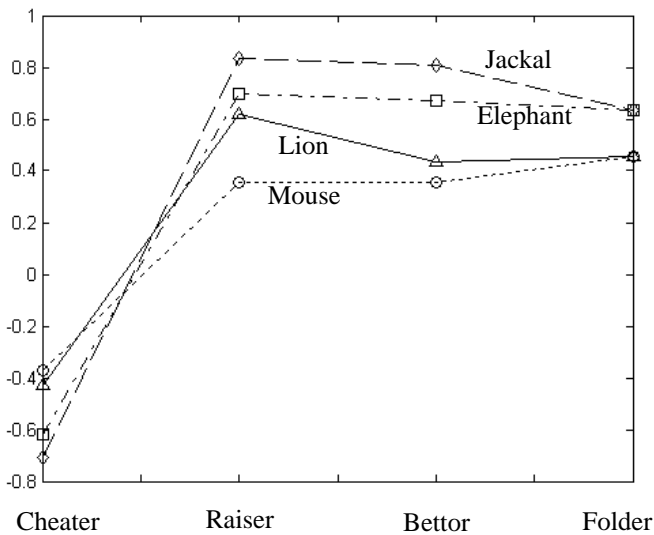
Figure 6. Results of pair-wise face-offs between each Novice style (Lion, Jackal, Elephant, Mouse) and each Extreme style (Folder, Bettor, Raiser, Cheater).

In Figure 6, positive numbers mean the Novice style plotted by a line (Lion, Jackal, Elephant, Mouse) wins against the Extreme style shown along the x-axis, while negative numbers mean the Novice style loses to the Extreme style. Here we see that the Loose styles (Jackal and Elephant), which are the worst styles in Novice face-offs (Figure 5), are the best styles in face-offs against the Extreme styles. However, against a Cheater, the Tight styles (Mouse and Lion) are the best.

Finally, referring to the parenthetic numbers in Figure 5, all of the Novice styles are remarkably effective against the Eagle who plays with Expert skill. The Mouse is the best against the Eagle (losing only 0.08 chips/game), and for a given style the best Novice opponent is almost as effective as the Expert opponent (e.g., against the Jackal, the Lion and the Expert both earn about 0.21 chips/game).

Discussion

The above results (Figures 5 and 6) illustrate three major findings. First, there are measurable differences between Novice styles playing against each other, with net earnings of 0.1-0.2 chips/game (Figure 5). Second, the Novice styles are all very effective against Extreme styles, with net earnings of 0.4-0.8 chips/game (Figure 6) – but the Mouse, who is most effective against other Novices (Figure 5), is least effective against the Extremes (Figure 6). Finally, Expert skill beats all Novice styles, but the winnings are in the same ballpark (0.1-0.2 chips/game) as the Novice styles against other Novice styles.

These findings demonstrate that the right style can be effective against other styles and against skill. The question is: How can the simple *style* of a Novice (e.g., Mouse) be so effective against the complex *skill* of the Expert (Eagle)? The answer, I think, involves two factors.

The first factor is that winning poker takes a lot of luck and not that much skill when the game is played with "low limits" on bets and raises. This is because, in low limit games, the player who is dealt the best hands can usually win the most chips simply by betting only when he has a decent hand. In fact, this is the reason that poker pros prefer "high limit" or even "no limit" games, where advanced strategies like bluffing become important – because a good bluff allows a player to win chips even when he does not have the best hand (Scarne 1980).

For example, consider node A' in Figure 3. In the low limit game analyzed here, the pot contains 8 chips when Player A is faced with the decision to call or fold (after Player B raises). With cost:pot stakes of 2:8=1:4, Player A should call at node A' with almost any hand that was worth betting in the first place (at node A). But now consider a no limit game where Player B raises Player A's initial bet. If Player B makes a small raise then Player A should call. But if B makes a large raise then the cost:pot stakes become such that A will need nearly even win:loss odds to justify a call. Furthermore, since B just made a large raise, A will probably think that B has a good hand so A will probably update (reduce) his win:loss odds in light of the raise. In this way, a bluff by B can get A to fold and B can win the pot even with a bad hand.

The second factor that accounts for the Mouse's success is that his simple strategy is functionally equivalent to a much more complex strategy (see Gigerenzer and Todd 1999). In particular, the Mouse's heuristic style (which requires a better than average card to bet and a significantly higher card to raise) implicitly captures two features of strategic skill, namely: (a) the projected cost:pot stakes are typically such that one should raise only when one has a relatively high chance of winning a showdown (high win:loss odds) or a relatively high chance of causing one's opponent to fold, and (b) one's chance of winning a showdown or chance of an opponent folding are typically reduced after an opponent has raised if the opponent is also considering (a).

The results (Figures 5 and 6) also demonstrate that the *best style* depends on an opponent's style. In particular, if one is facing an Expert or Novice of unknown style then the best style is a Mouse (Tight and Passive). But if one is facing a known Jackal then the best style is a Lion (Tight and Aggressive). And if one is facing an Extreme style then the best style is a Jackal (Loose and Aggressive). These results can be summarized in two rules of thumb that correspond to two colloquial adages: *Tight is right* (for the Mouse and Lion) and *Fight fire with fire* (for the Lion and the Jackal).

It is interesting that the most effective Novice strategies in *Pared-down Poker* are similar to the simplistic strategies that have been found to be effective in other mind games. For example, in his seminal work on the cooperative-competitive game known as the Prisoner's Dilemma, Robert Axelrod (1997; 1984) found that the best strategy was *Tit for tat*, which involves prompt and consistent response to an opponent's aggression. This strategy is Lion-like in that it combines the best of the Mouse and the Jackal. That is, in *Tit for tat* one plays cooperatively by the rule *Tight is right* unless the other player gets aggressive, in which case one plays competitively by the rule *Fight fire with fire*.

Conclusion

This paper developed formal models of *style* and *skill* in poker playing. Here, style is defined as a *functional* basis for making decisions in the absence of a *rational* basis for making decisions, and different styles are assumed to stem from mental limits of *precision* and *projection* in estimating odds and stakes. Using a binary distinction between *pessimistic* and *optimistic,* Novice styles were characterized and compared to Expert skill.

The computational models for style and skill were used to answer specific questions in the context of a simplified poker game, such as: Which Novice style is best, and how good are Novice styles against Expert skill? The results show that the best style depends on an opponent's style, and that the right style can be remarkably effective against other styles and skill.

Both the general approach and the specific results of this study have implications beyond poker playing. The general approach applies to any domain where the basis for choice (i.e., expected utility) is clouded by mental limits and hence where choices are governed by style as well as skill. The approach is to analyze quantities on which decisions are or should be based (e.g., odds and stakes) and then characterize the cognitive styles (e.g., Tight or Loose) by which these quantities are estimated by human minds. In this way, researchers can answer questions like: What are the various styles? Why are these styles adopted? When do various styles perform well against (or along with) each other and against (or along with) skill? How can styles be improved (by training and/or systems) so that people can perform at higher levels of skill?

The specific results of this study extend beyond poker playing because similar challenges of *inference* and *investment* arise in real world domains like business and warfare. The results of this study show there is a computational basis for stylistic heuristics, and the details of this study shed light on when a strategy like *Tight is right* works well and when it breaks down such that another strategy like *Fight fire with fire* becomes better.

References

Axelrod, R. 1997. *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration.* Princeton, NJ: Princeton University Press.

Axelrod, R. 1984. *The Evolution of Cooperation.* Basic Books.

Billings, D., Burch, N., Davidson, A., Holte, R., Schaeffer, J., Schauenberg, T., and Szafron, D. 2003. Approximating Game-Theoretic Optimal Strategies for Full-Scale Poker. *Proceedings of the International Joint Conference on Artificial Intelligence.*

Billings, D., Aaron, D., Schaeffer, J., and Szafron, D. 2002. The Challenge of Poker. *Artificial Intelligence Journal* 134:201-240.

Epstein, R. A. 1977. *The Theory of Gambling and Statistical Logic.* San Diego: Academic Press.

Gigerenzer, G., and Todd, P. M. 1999. *Simple Heuristics That Make Us Smart.* New York: Oxford University Press.

Hellmuth, P. 2003. *Play Poker Like the Pros.* New York: Harper-Collins.

Koller, D., and Pfeffer, A. 1997. Representations and Solutions for Game-Theoretic Problems. *Artificial Intelligence* 94:167-215.

Kuhn, H. W. 1950. A Simplified Two-Person Poker. *Contributions to the Theory of Games, Vol. 1.* Princeton, NJ: Princeton University Press, 97-103.

McDonald, J. 1950. *Strategy in Poker, Business and War.* New York: Norton.

McManus, J. 2003. *Positively Fifth Street: Murderers, Cheetahs and Binion's World Series of Poker.* New York: Farrar, Straus and Giroux.

Nash, J. F., and Shapley, L. S. 1950. A Simple Three-Person Poker Game. *Contributions to the Theory of Games, Vol. 1.* Princeton, NJ: Princeton University Press, 105-116.

Scarne, J. 1980. *Scarne's Guide to Modern Poker.* New York: Simon and Schuster.

Schoonmaker, A. N. 2000. *The Psychology of Poker.* Las Vegas: Two Plus Two Publishing.

Simon, H. 1981. *The Sciences of the Artificial.* Cambridge, MA: MIT Press.

Sklansky, D. 1987. *The Theory of Poker.* Las Vegas: Two Plus Two Publishing.